

REMARKS

The Office Action dated June 14, 2004 has been received and carefully noted. The above amendments to the claims, and the following remarks, are submitted as a full and complete response thereto.

Claims 1-4, 6, 8, 10, 12, 13, 15, 16, 17 and 32 have been amended. Claim 27 has been cancelled. No new matter has been added, and no new issues are raised which require further consideration and/or search. Claims 1-15, 17-26, 28, 30-43 and 45 are submitted for consideration.

Claim 2 was rejected under 35 U.S.C. 102(b) as being anticipated by U.S. Patent No. 5,463,620 to Sriram. The rejection is traversed as being based on a reference that neither teaches nor suggests the novel combination of features clearly recited in independent claim 2.

Claim 2 recites a method including dividing a total amount of data, based upon an individual weight assigned to each of a plurality of queues, into an amount of data that each of said queues may service. The method also includes scheduling packets in the plurality of queues during scheduling cycles. Each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions. The method further includes servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles. A number of scheduling

cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth. The method also include servicing one or more populated queues, each of said servicing of a populated queue continuing until said populated queue is no longer populated or said amount of data determined for said populated queue has been released; and servicing one or more of said queues that remain populated, if any, until said total amount of data has been released from all of said queues in combination including said servicing of said populated queues.

As will be discussed below, the cited prior art reference of Sriram fails to disclose or suggest the elements of any of the presently pending claims.

Sriram teaches an ATM communications network that includes a plurality of interconnected nodes. The nodes include input and output links and each output link from a node is provided with a queuing mechanism which receives ATM cells for delivery into a channel on the output link. Traffic which is to be directed on the output link is selectively directed to a number of different queuing circuits based on the results of classification. The traffic in the queuing circuits is guaranteed a minimum percentage of the total bandwidth available on the output link. A server defines a cycle time period during which it will retrieve cells from all of the queues having cells to send. The server divides the cycle time period into slices, assigns a time slice to each queue and permits

each queue to empty cells onto the output link during its respective time slice. The server visits each queue in sequence and removes a predetermined number of cells from each queue. All queues are visited within the next cycle time defined by the server. At the completion of the cycle time period, the server repeats the cycle of visiting each queue and removing respective predetermine number of cells from the queue. If any one of the queues contains no cells, then the server completely passes over the empty queue and moves on to the next queue in the sequence. If any of the queues contain a number of cells which is less than the predetermined number of cells the server is scheduled to remove during the cycle time period, then the server removes cells from that queue until it is empty and moves on to the next queue to remove its allotted number of cells. Col. 5, line 51- Col. 6, line 46.

Applicant submits that Sriram does not teach or suggest each of the features now clearly recited in claim 2. Sriram simply does not teach or suggest partitioning each scheduling cycle into regions, each of which is coextensive with a highest bandwidth being managed by a node and each scheduling cycle is coextensive with highest counting modulo partitions as recited in claim 2. Sriram also does not teach or suggest servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth

queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth as recited in claim 2. Instead, Sriram teaches that a server defines a cycle time period during which it will retrieve cells from all of the queues having cells to send and the server divides the cycle time period into slices, assigns a time slice to each queue and permits each queue to empty cells onto the output link during its respective time slice. The present invention, on the other hand, allows for partitioning each scheduling cycle into regions where each region support the highest bandwidth being managed by a node and each scheduling cycle support highest counting modulo partitions. According to the present invention, after each scheduling cycle is defined, queues associated with the highest bandwidth in are serviced in at least one partition during each scheduling cycle and consecutive bandwidth partitions of queues associated with lower bandwidths are serviced across several cycles. The present invention thus ensures fairness by ensuring that low bandwidth queues are processed in each queuing cycle even though consecutive bandwidth partitions of a queue associated with a lower bandwidth may not be processed in consecutive cycles. Therefore, Applicant respectfully asserts that the rejection under 35 U.S.C. §102(b) should be withdrawn.

Claims 1, 3-14, 16, 17 19 and 32 were rejected under 35 U.S.C. 102(b) as being unpatentable over Sriram in view of U.S. Patent No. 6,094,435 to Hoffman et al. The rejection is traversed as being based on references that neither teach nor suggest the novel combination of features clearly recited in independent claims 1, 3, 17 and 32.

Claim 1 recites an apparatus including a plurality of users and resources that are partitioned according to a ranking of bandwidth associated with users. The resources are partitioned according to a highest bandwidth supported by a node and an amount of bandwidth provided to each of the plurality of users is ranked from highest to lowest. The apparatus also includes a queue scheduler that a) schedules one or more packets within the node during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with the highest bandwidth supported by the node and each schedule cycle is coextensive with a highest counting modulo partitions, and b) services users associated with the highest bandwidth in at least one partition during each scheduling cycle and services consecutive bandwidth partitions of user associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the user decreases and the partition spacing for servicing a lower bandwidth user is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth.

Claim 3 , upon which claims 4-16, 18-26, 28-31, 33-43 and 45, depend recites a method including distributing a partition worth of data across a plurality of queues according to a weight assigned to each of said queues so that each of said queues has its own sub-partition worth of data, each of the queues being capable of holding one or more packet identifiers. Each of the one or more packet identifiers points to its own packet and plurality of queues ranges from a highest priority queue to a lowest priority queue. The

method further includes scheduling packets in the plurality of queues during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions. The method also includes servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth. The method includes yet still flowing a flow of one or more packet identifiers from an active populated queue, until: its unpopulated if less than its sub-partition worth of data has flowed, or until its sub-partition worth of data has flowed, or until the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole.

Claim 17 recites a method including distributing a partition worth of data across a plurality of queues according to a weight assigned to each of said queues so that each of said queues has its own sub-partition worth of data, each of the queues being capable of holding one or more packet identifiers. Each of the one or more packet identifiers points to its own packet and plurality of queues ranges from a highest priority queue to a lowest

priority queue. The method further includes scheduling packets in the plurality of queues during scheduling cycles, wherein each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions. The method also includes servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles, wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth. The method includes yet still flowing a flow of one or more packet identifiers from an active populated queue, until: its unpopulated if less than its sub-partition worth of data has flowed, or until its sub-partition worth of data has flowed, or until the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole. A populated queue is deemed active if it is the highest priority populated queue out of those of the populated queues that have not yet been deemed active, such that populated queues are deemed active in succession until the lowest priority populated queue has been deemed active or until the combination of flows from

those of the queues that have been active results in the partition worth of data having flowed from the those of the queues that have been active, as a whole.

Claim 32 recites an apparatus including a scheduler that schedules packets in the plurality of queues during scheduling cycles. Each scheduling cycle is partitioned into regions that are coextensive with a highest bandwidth being managed by a node and each schedule cycle is coextensive with a highest counting modulo partitions. The scheduler also services queues associated with the highest bandwidth in at least one partition during each scheduling cycle and services consecutive bandwidth partitions of queues associated with lower bandwidths across several cycles. The number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases. Partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth; and. The scheduler controls a flow of one or more packet identifiers from an active populated queue, until 1) its unpopulated if less than its sub-partition worth of data has flowed, or until 2) its sub-partition worth of data has flowed, or until 3) the combination of flows from those of said queues that have been active results in said partition worth of data having flowed from said those of said queues that have been active, as a whole.

As will be discussed below, the combination of Sriram and Hoffman fails to disclose or suggest the elements of any of the presently pending claims.

Hoffman teaches a system which includes a multi-layer network element, various networks, end stations, router and bridges. The multi-layer network element includes a processor, various memory locations, a switching elements and a plurality of network element ports. Col. 8, lines 38-49. The switching element includes input ports, a forwarding logic, a packet memory manager, and output ports. Col. 9, lines 9-15. Packets in the network element are buffered at each output port before the packet is transmitted across the physical medium to the next or final destination. Queuing at the input and output ports are based on pointers. Col. 18, lines 35-39. Each output port has a plurality of output queues and each queue has a pair of pointer registers to indicate the beginning and the end of the queue. Col. 18, lines 49-67. The forwarding logic passes global priority information to the input and output ports to classify packets into different queues. The output port uses the global priority information to determine to which queue a given packet will be forwarded. Col. 19, lines 28-40. The output port also includes a scheduler which allocates fixed rates to each queue for transmission within the output port. Each queue in each output port has associated with it three programmable registers which contain the weights to be used for their associated queue. In one embodiment, the scheduler will not service a lower priority queue as long as there are packets in the higher priority queues. In another embodiment, the scheduler polls each queue and services the packets based on the associated weight of the queue before servicing the next queue. In yet another embodiment, the scheduler attempts to enforce the rates over several polling rounds that comprise a frame. Col. 20, lines 24-67. The scheduler services a queue and

decrements a transmit register according to the number of bytes transferred until the value of the transmit register is equal to or less than zero. Then the scheduler starts processing the next queue in the round and updates the transmit register of the just serviced queue by adding to the transmit register a quantum of bytes as represented by the value in the weight register. As such, a queue may finish transmitting a packet even if the number of bytes to finish transmitting the packet causes the value of the transmit register to drop below zero. This allows the scheduler to take into account for the queue in the subsequent round or frame, any overrun in the current round or frame. Therefore, when the value in the weight register is added to the value in the transmit register, the number of packets that the queue may transmit during the next round, or frame is reduced by the amount the queue went over its allocation for the current round or frame. Col. 21, lines 17-35.

Applicant submits that combination of Sriram and Hoffman does not teach or suggest each of the features now clearly recited in independent claims 1, 3, 17 and 32 and the claims dependent thereon. As explained above in the arguments relating to claim 2, Sriram simply does not teach or suggest partitioning each scheduling cycle into regions, each of which is coextensive with a highest bandwidth being managed by a node and each scheduling cycle is coextensive with highest counting modulo partitions as recited in claims 1, 3, 17 and 32. Sriram also does not teach or suggest servicing queues associated with the highest bandwidth in at least one partition during each scheduling cycle and servicing consecutive bandwidth partitions of queues associated with lower bandwidths

across several cycles wherein a number of scheduling cycles between servicing of consecutive bandwidth partitions increases as the bandwidth associated with the queue decreases and partition spacing for servicing a lower bandwidth queue is determined by multiplying a number of lower bandwidth users that can be serviced by the next highest bandwidth by a partition modulo of the next highest bandwidth as recited in claims 1, 3, 17 and 32.

Hoffman does not cure the deficiencies of Sriram outlined above. Instead Hoffman teaches that the scheduler will not service a lower priority queue as long as there are packets in the higher priority queues. Hoffman also teaches that in an embodiment, the scheduler polls each queue and services the packets based on the associated weight of the queue before servicing the next queue. Hoffman further teaches that the scheduler services a queue and decrements a transmit register according to the number of bytes transferred until the value of the transmit register is equal to or less than zero. Then the scheduler starts processing the next queue in the round and updates the transmit register of the just serviced queue by adding to the transmit register a quantum of bytes as represented by the value in the weight register. As such, a queue may finish transmitting a packet even if the number of bytes to finish transmitting the packet causes the value of the transmit register to drop below zero. Applicant submits that the teaching of Hoffman is quite different from the claimed invention. Furthermore the combination of Hoffman and Sriram also does not teach the scheduling of the claimed invention. Therefore, Applicant asserts that the rejection under 102(b) should be withdrawn because

neither Sriram nor Hoffman, whether taken singly or combined, teaches or suggests each feature of now clearly recited in claim 1, 3, 17, and 32 and hence, dependent claim 4-14, 16, and 19 thereon.

Claims 15, 18, 20-28, 30, 31, 33-43 and 45 were rejected under 35 U.S.C. 103(a) as being unpatentable over Sriram in view of Hoffman et al. as applied to claim 1 above, and further in view of U.S. patent No. 5,982,748 to Yin et al. The rejection is traversed as being based on references that neither teach nor suggest the features in independent claim 3, upon which claims 15, 18, 20-26, 28, 30, 31, 33-43 and 45 depend.

Yin et al. teaches a system for controlling admission of connection requests and allocating bandwidth to ensure efficient use of network resources. According to Yin et al., an allocation factor indicates whether an associated service class is fully booked, under-subscribed or over subscribed. A service class is fully booked if the maximum allowable subscribed bandwidth equals to the bandwidth allocated to the service class. The service class is over-subscribed if the maximum allowable subscribed bandwidth is greater than the bandwidth allocated to the service class. The service class is under subscribed if the maximum allowable subscribed bandwidth less than the bandwidth allocated to the service class.

Yin et al. does not cure the deficiencies in Sriram and Hoffman as to claim 3 as explained above. Therefore, Applicant respectfully asserts that the rejection under 35 U.S.C. §103(a) should be withdrawn because neither Sriram, Yin et al. nor Hoffman, whether taken singly or combined, teaches or suggests each feature of claim 3 and hence

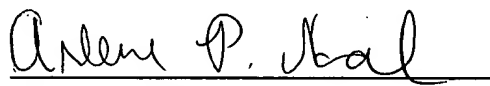
dependent claims 15, 18, 20-26, 28, 30, 31, 33-43 and 45.

As noted previously, claims 1-15, 17-26, 28, 30-43 and 45 recite subject matter which is neither disclosed nor suggested in the prior art references cited in the Office Action. It is therefore respectfully requested that all of claims 1-15, 17-28, 30-43 and 45 be allowed and this application passed to issue.

If for any reason the Examiner determines that the application is not now in condition for allowance, it is respectfully requested that the Examiner contact, by telephone, the applicant's undersigned attorney at the indicated telephone number to arrange for an interview to expedite the disposition of this application.

In the event this paper is not being timely filed, the applicant respectfully petitions for an appropriate extension of time. Any fees for such an extension together with any additional fees may be charged to Counsel's Deposit Account 50-2222.

Respectfully submitted,



Arlene P. Neal
Registration No. 43,828

Customer No. 32294
SQUIRE, SANDERS & DEMPSEY LLP
14TH Floor
8000 Towers Crescent Drive
Tysons Corner, Virginia 22182-2700
Telephone: 703-720-7800
Fax: 703-720-7802

APN:lls